Basic Multicast Troubleshooting

Piotr Wojciechowski (CCIE #25543)

About me

- Senior Network Engineer MSO at VeriFone Inc.
- Previously Network Solutions Architect at one of top polish IT integrators
- CCIE #25543 (Routing & Switching)
- Blogger – http://ccieplayground.wordpress.com
- Administrator of CCIE.PL board
  - The biggest Cisco community in Europe
  - Over 6100 users
  - 3 admin, 7 moderators
  - 48 polish CCIEs as members, 20 of them actively posting
  - About 150 new topics per month
  - About 1000 posts per month
  - English section available!

AGENDA

1. Multicasts – basics you have to know
2. Basic diagnostics
3. Why can't I register in multicast network?

Multicasts – basics you have to know

- Multicasts:
  - UDP Traffic
  - Best effort
  - State of multicast distribution and signalization changes dynamically
  - Application controlled
  - Multicast replication
  - Both source and receiver have to register

Multicasts – basics you have to know

- One source – many receivers
  - Source have unicast A, B or C class address
  - Receiver is unknown, destination address is from class D

MULTICASTS – BASICS YOU HAVE TO KNOW

- Three modes:
  - ***Dense Mode*** – we assume that receiver is on every subnet, so multicast traffic should be replicated on every network. Routers must explicitly declare they do now want to receive multicast traffic for particular group (due to no receivers) and refresh this request periodically.
  - *Sparse Mode* – unidirectional tree is built from defined point in network (RP) toward receivers, that have explicitly declare they want to receive traffic. In next stages this tree can be rebuild using source of multicast traffic as a root.

Multicasts – basics you have to know

- ***Sparse-Dense Mode***
  - Hybrid mode. Groups with defined RP use sparse mode. All others use dense mode
  - This mode is required to propagate information of RP if Auto-RP technology is used but without Auto-RP Listener.

Multicasts – basics you have to know

- Building multicast tree:
  - *Shared Tree* – root of the tree is RP. All receivers have to register to RP, he manage traffic distribution to receivers.

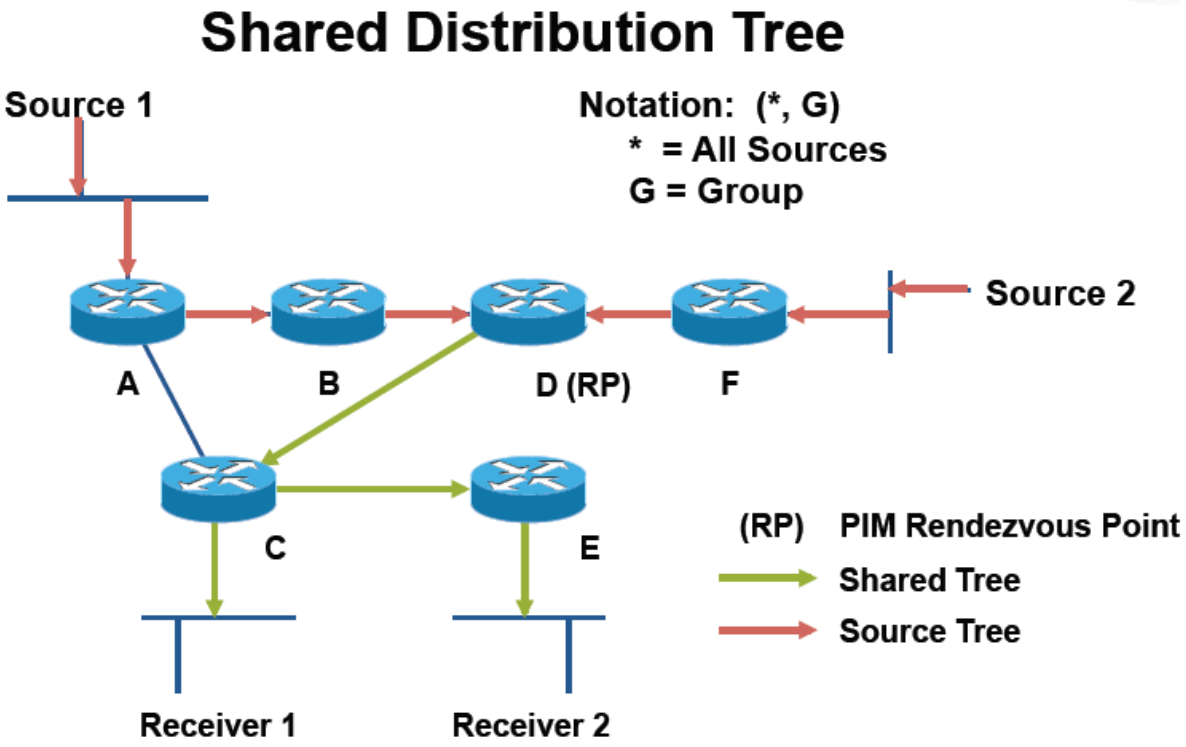

Multicasts – basics you have to know

- Building multicast tree:
  - *Source Tree* – root of the tree is source for prticular multicast group. For each receiver (S,G) pair is created.

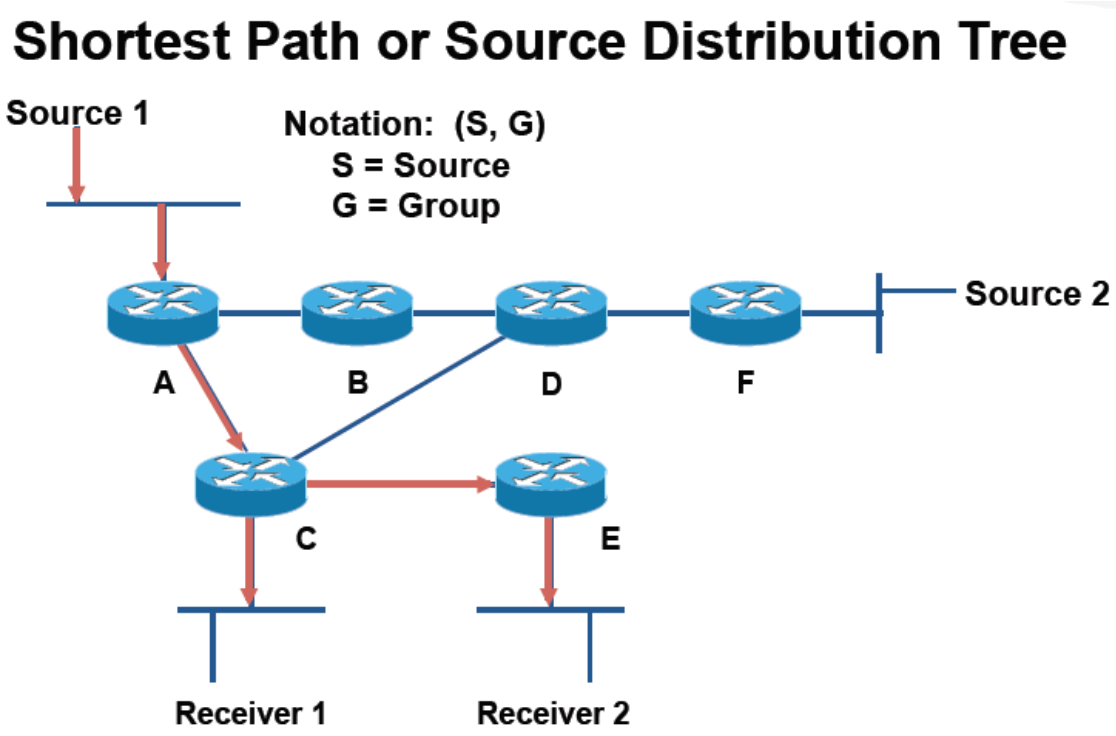

Multicasts – basics you have to know

- Any Source Multicast (ASM)
  - Classic mode of PIM-SM
  - Shared Tree and Source Tree are used
  - For Shared Tree RP's are used
- Source Specific Multicast (SSM)
  - Only Source Tree is used
  - No RP
  - Multicast groups only within `232.0.0.0/8` (IPv4) and `FF3x::/96` (IPv6) subnets
- Bidirectional PIM (BiDir)
  - Shared Tree

Multicasts – basics you have to know

- (*,G) – Shared Tree notification
  - RP is a root
  - We know group address but source is unknown
- (S,G) – Source Tree notification
  - Source of multicast traffic is a root
  - We know both source and group address
- IIF – Incoming Interface
  - Interface towards multicast source (Source Tree) or RP (Shared Tree)
- OIL – Outgoing Interface List
  - List of interfaces where multicasts are replicated and distributed

Multicasts – basics you have to know

- RP – *Randezvous Point*. For Sparse Mode point of reference required to build common multicast tree.
- FHR – *First Hop Router*. First router on the path. Responsible for source registration at RP.
- LHR – *Last Hop Router*. Last router on the path, closes to the receiver.

Multicasts – basics you have to know

- Reverse Path Forwarding Check
  - Used to avoid loops while distributing multicasts
  - Source address is used as a reference:
    - *If best route in unicast routing table to source is through the interface, where the packet was received, then is should be replicated*
    - *If packed is received with other interface then it should be dropped*
  - If same packet is received through multiple interfaces it will be replicated only once.

Reverse Path Forwarding Check

- Both packets with data (data plane) as well as some of control packets (control plane) have to pass RPF rule:
  - PIM (*,G) Join – are always sent using shortest path to RP
  - BSR/RP addresses sent inside of BSR messages
  - Every multicast packet in data plane

Reverse Path Forwarding Check

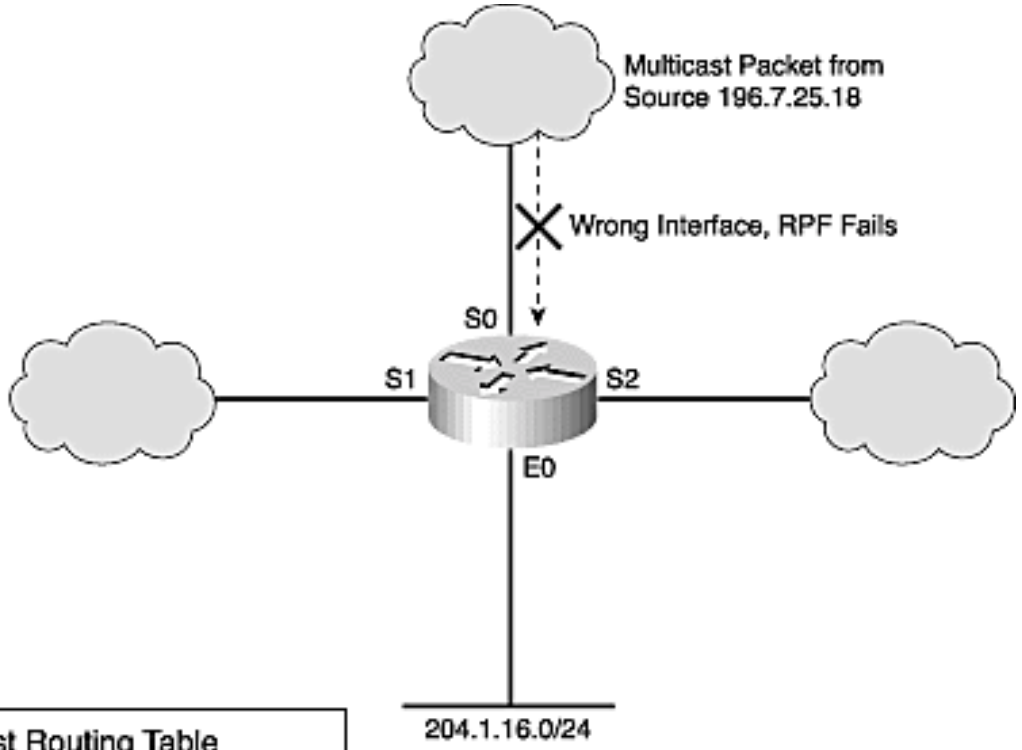

Unicast Routing Table	
Network	Interface
196.7.25.0/24	S1
196.7.26.0/24	S0
204.1.16.0/24	E0

Packet arrived on wrong interface, therefore packet should be dropped.

Reverse Path Forwarding Check

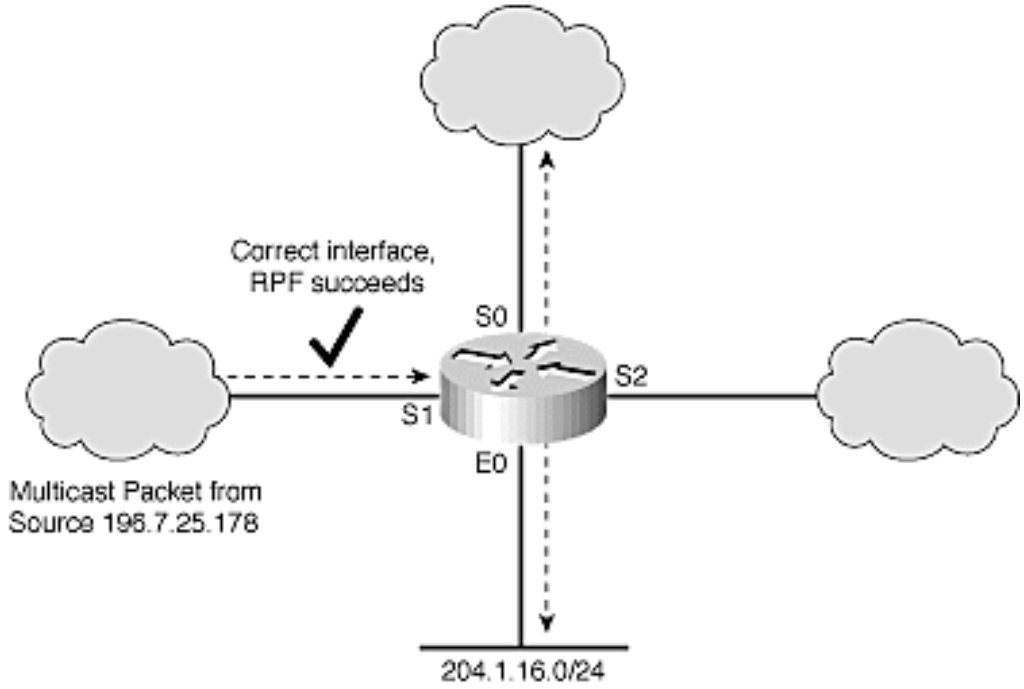

Unicast Routing Table	
Network	Interface
196.7.25.0/24	S1
196.7.26.0/24	S0
204.1.16.0/24	E0

Packet arrived on correct interface, therefore packet should be replicated.

Basic commands

```
R8#show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.100.0.1), 00:00:33/stopped, RP 10.10.0.1, flags: SJCL
  Incoming interface: Ethernet0/2, RPF nbr 10.0.38.3
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse-Dense, 00:00:32/00:02:56

(10.10.0.1, 224.100.0.1), 00:00:30/00:02:29, flags: LJT
  Incoming interface: Ethernet0/2, RPF nbr 10.0.38.3
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse-Dense, 00:00:30/00:02:56
```

Flags

(*,G)
Shared Tree

(S,G)
Source Tree

Basic commands

```
R8#sh ip igmp interface
Ethernet0/0 is up, line protocol is up
  Internet address is 10.0.89.8/24
  IGMP is enabled on interface
  Current IGMP host version is 2
  Current IGMP router version is 2
  IGMP query interval is 60 seconds
  IGMP configured query interval is 60 seconds
  IGMP querier timeout is 120 seconds
  IGMP configured querier timeout is 120 seconds
  IGMP max query response time is 10 seconds
  Last member query count is 2
  Last member query response interval is 1000 ms
  Inbound IGMP access group is not set
  IGMP activity: 1 joins, 0 leaves
  Multicast routing is enabled on interface
  Multicast TTL threshold is 0
  Multicast designated router (DR) is 10.0.89.8 (this system)
  IGMP querying router is 10.0.89.8 (this system)
  Multicast groups joined by this system (number of users):
      224.100.0.1(1)
```

Basic commands

```
R8#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter   Group Accounted
224.100.0.1      Ethernet0/0              00:07:19  00:02:56  10.0.89.8
224.0.1.40       Ethernet0/1              00:30:09  00:02:47  10.0.78.8
```

Basic commands

```
R8#sh ip igmp groups  detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source,
       Ac - Group accounted towards access control limit

Interface:  Ethernet0/0
Group:                  224.100.0.1
Flags:                  L U
Uptime:                 00:07:30
Group mode: EXCLUDE (Expires: 00:02:45)
Last reporter:          10.0.89.8
Source list is empty

Interface:  Ethernet0/1
Group:                  224.0.1.40
Flags:                  L U
Uptime:                 00:30:20
Group mode: EXCLUDE (Expires: 00:02:36)
Last reporter:          10.0.78.8
Source list is empty
```

Sample Network

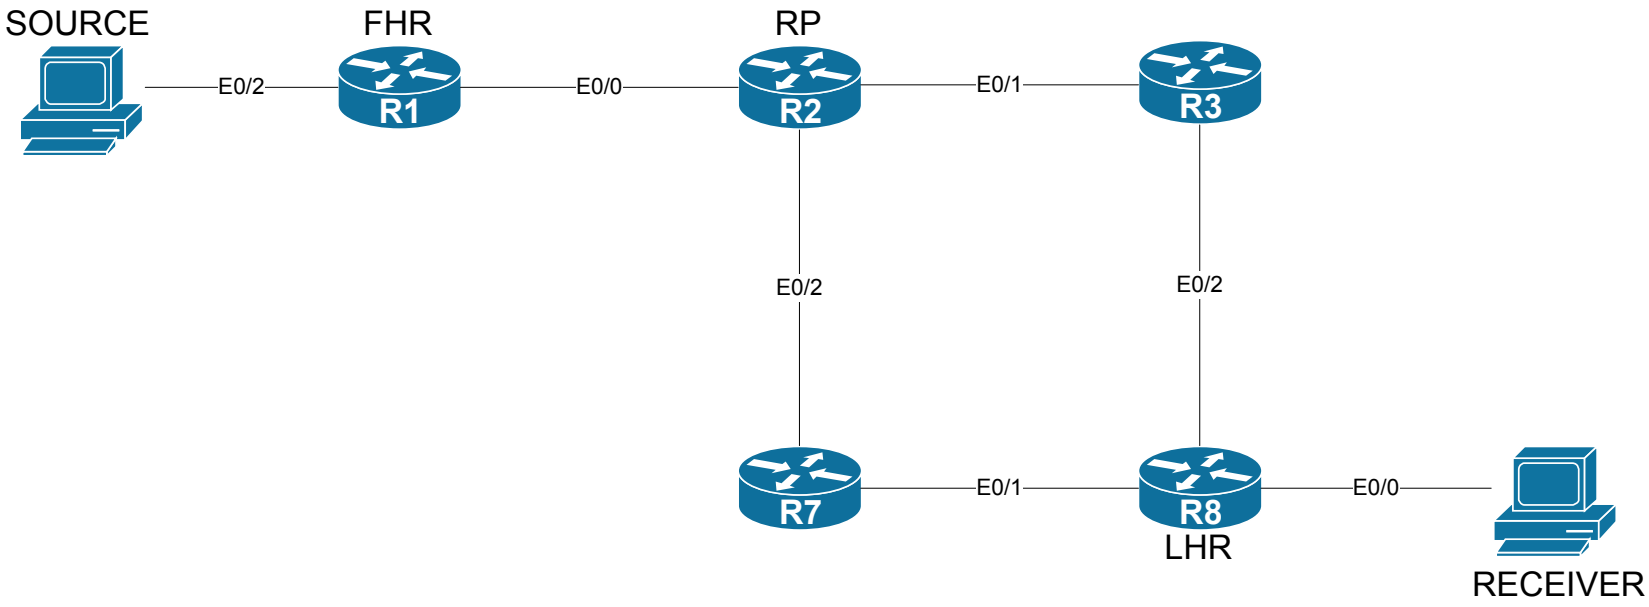

When we can expect problems?

- Many IGP protocols and redistribution between them – possible routing asymmetry and many topology changes
- PIM is not activated on all interfaces
- NBMA segments
- Many tunnels used to transport multicasts
- Problems with hardware performance
- Bad network design and RP placement

Troubleshooting approach

- Debugging from receiver to the source
  - Problem can affect only group of receivers – look for their common attributes
- Perform full diagnostics on every router on the path

PROBLEM: (S,G) TREE IS NOT CREATING

- Problem: in multicast routing table (mroute) on LHR we can only see (*,G) entries
  - We do have active source and its sending traffic
  - (*,G) path is correct
  - No (S,G) entries in mroute table

PROBLEM: (S,G) TREE IS NOT CREATING

- Solution:
  - We need a receiver!
    - Source-specific multicast (SSM) tree would not build unless there is at least one receiver for a group
    - Properly registered receiver is the one, that sent PIN Join for multicast group

PROBLEM: DUPLICATED PACKETS

- Problem: Receiver sometimes receive duplicated packets for groups he joined to.
- Possible causes: Traffic sent in dense mode
  - Router periodically floods data on interfaces time to time
  - If no receivers are registered then multicast traffic is pruned unless timer expire
  - If there are at least two routers on network segment they elect forwarder who will replicate traffic on this segment
  - When the timers expires process of flooding, pruning and election are repeated which may cause of duplicated packets being sent over that segment

PROBLEM: DUPLICATED PACKETS

- Solution:
  - Change dense mode to sparse mode
  - This may also be caused by hardware problems, ie. overloaded line cards in routers or switches.
  - It's up to an application to handle duplicated packets, so don't worry to much (unless it kills your bandwidth)

PIM-SM REGISTRATION PROCESS

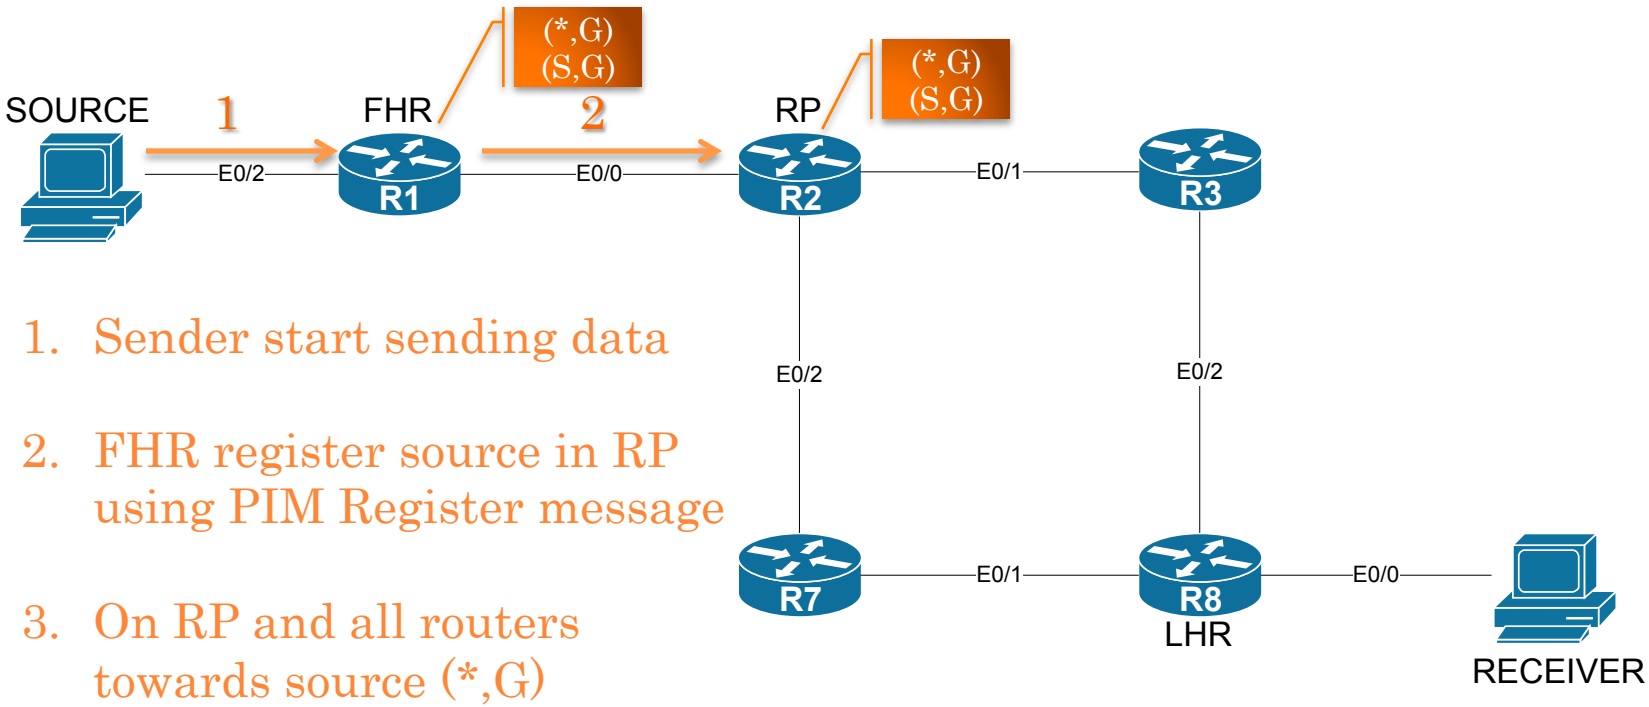

1. Sender start sending data
2. FHR register source in RP using PIM Register message
3. On RP and all routers towards source (*,G) and (S,G) entries are created

PIM-SM REGISTRATION PROCESS

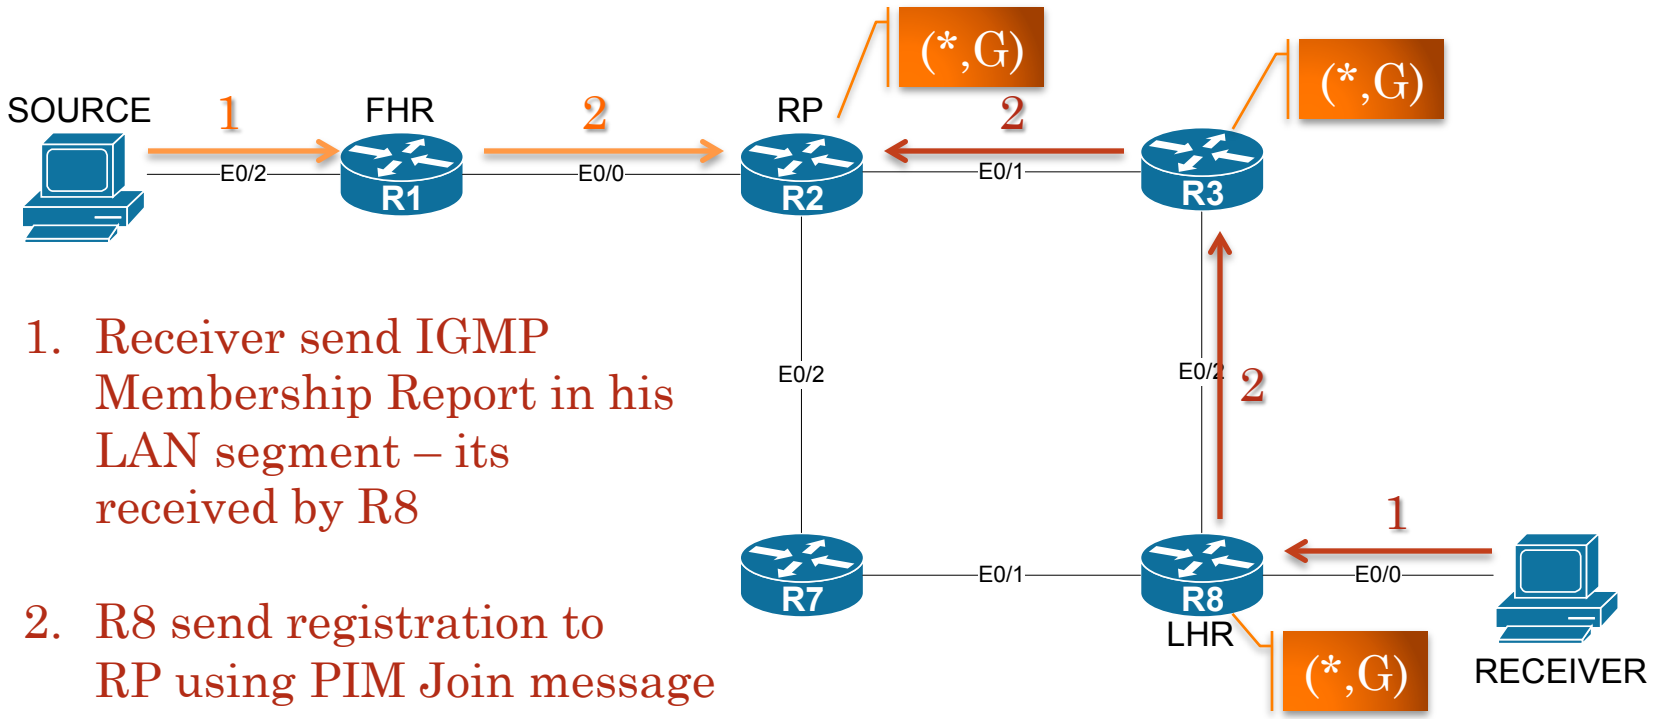

1. Receiver send IGMP Membership Report in his LAN segment – its received by R8
2. R8 send registration to RP using PIM Join message
3. (*,G) entries are created on all router on path to RP

PIM-SM REGISTRATION PROCESS

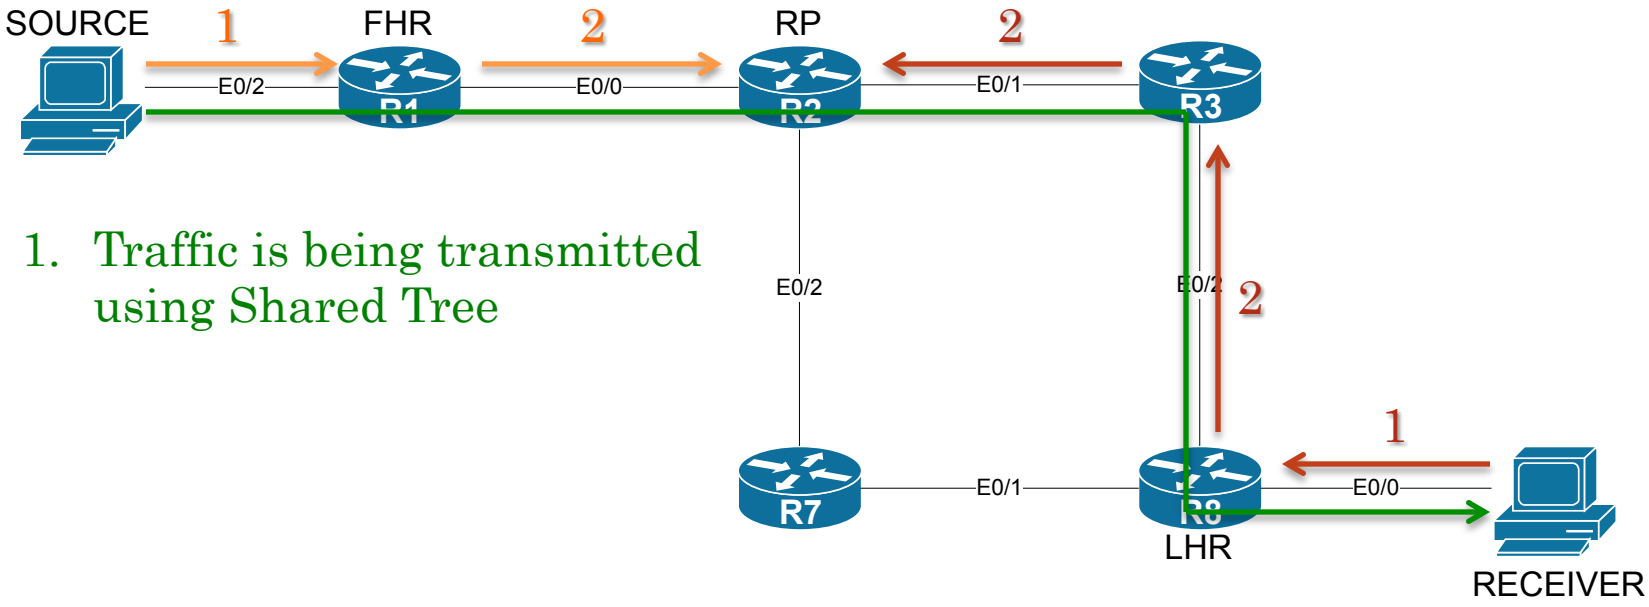

PROBLEM: RECEIVER CAN'T REGISTER

- **Problem 1:**
  - We are sure that receiver sends IGMP Membership Report
  - Registration request is not processed by LHR

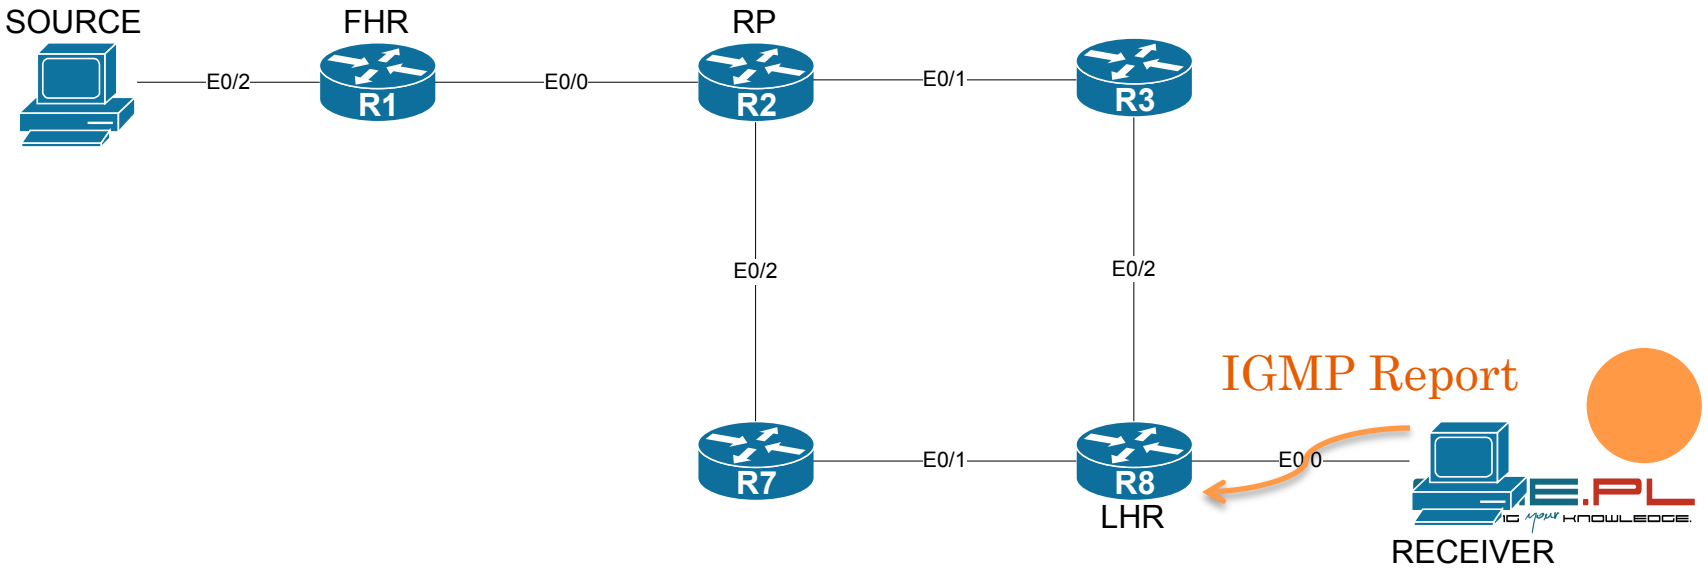

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes:
  - On L2 path from receiver to source IGMP filtering occurs
  - Access-list on router interface
  - PIM is not enabled on interface, where receiver is connected

```
R8#sh run int e0/0
Building configuration...

Current configuration : 83 bytes
!
interface Ethernet0/0
 ip address 10.0.89.8 255.255.255.0
 ip ospf 1 area 0
end
```

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes:
  - PIM is not enabled on interface, where receiver is connected

```
Ethernet0/0 is up, line protocol is up
  Internet address is 10.0.89.8/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes2
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing access list is not set
  Inbound  access list is not set
```

No registration for the group (224.10.0.1)

```
R8#sh ip igmp groups
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter
Group Accounted
224.0.1.40       Ethernet0/1              02:13:52  00:02:03  10.0.78.8
```

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes:
  - PIM is not enabled on interface, where receiver is connected

```
R8#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R8(config)#int e0/0
R8(config-if)#ip pim sparse-mode

R8#sh ip igmp groups
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter
Group Accounted
224.10.0.1       Ethernet0/0              00:00:03  00:02:56  10.0.89.8
224.0.1.40       Ethernet0/1              03:21:35  00:02:23  10.0.78.8
```

PROBLEM: RECEIVER CAN'T REGISTER

- **Problem 2:**
  - LHR router is receiving registration requests
  - PIM Join is not transmitted over the network
- Possible causes (1):
  - Routers are not forwarding IGMP Membership Report – no known RP for the group

```
R8#show ip pim rp 224.10.0.1
Group: 224.10.0.1, RP: 10.10.0.1, v2, uptime 03:30:41, expires 00:02:19

R8#show ip mroute
IP Multicast Routing Table

(*, 224.10.0.1), 00:21:06/00:02:30, RP 10.10.0.1, flags: SJCL
  Incoming interface: Ethernet0/2, RPF nbr 10.0.38.3
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse, 00:21:06/00:02:30
```

PROBLEM: RECEIVER CAN'T REGISTER

- **Problem 2:**
  - LHR router is receiving registration requests
  - PIM Join is not transmitted over the network

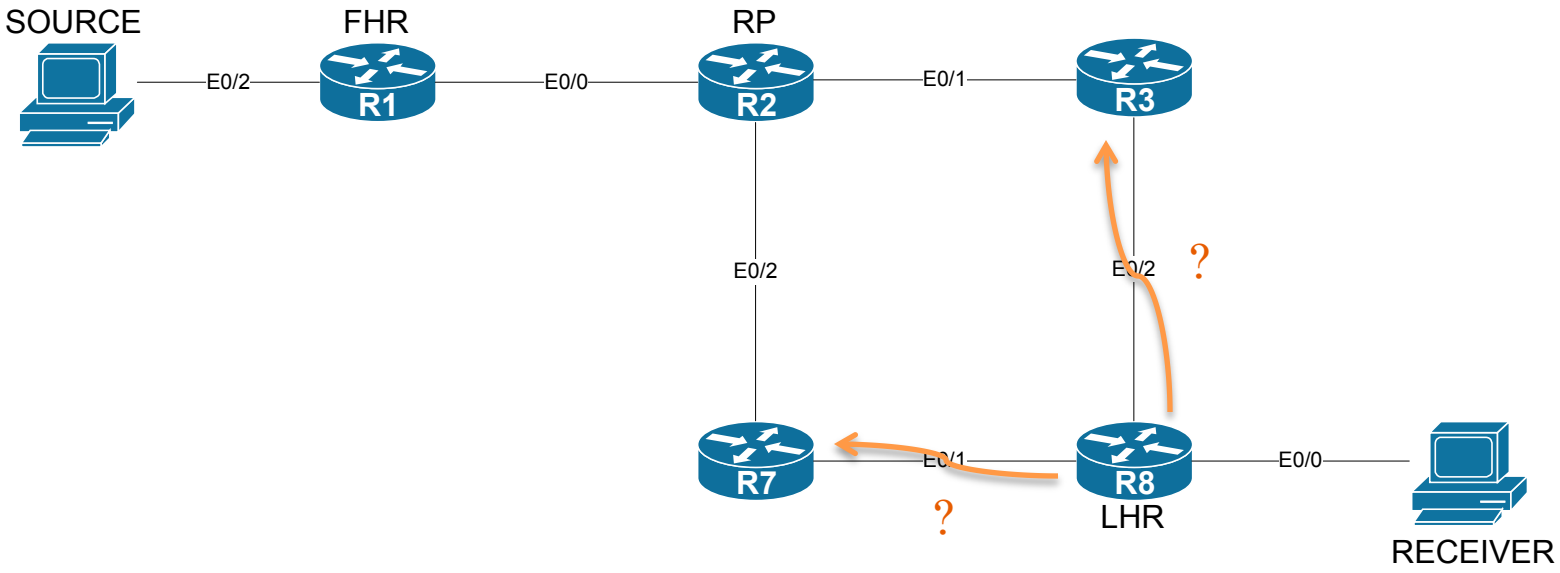

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (1):
  - PIM Join is not transmitted by router – no knowledge where RP is located

```
R8#sh ip pim rp mapping 224.10.0.1
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 10.10.0.1 (?), v2
    Info source: 10.10.0.1 (?), via bootstrap, priority 0, holdtime 150
         Uptime: 03:49:01, expires: 00:02:03
```

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (1):
  - PIM Join is not transmitted by router – no knowledge where RP is located
  - How router can know RP address?
    - Auto-RP
    - BSR
    - Static

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (2):
  - Router can't send PIM Join – no path to RP

```
R8#show ip mroute

(*, 224.10.0.1), 00:33:17/00:02:36, RP 192.168.2.2, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse, 00:33:17/00:02:36

R8#show ip rpf 192.168.2.2
 failed, no route exists

R8#show ip route 192.168.2.2
% Network not in table
```

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (3):
  - Router can't send PIM Join – incorrect path to RP

```
R8#show ip pim rp mapping 224.10.0.1
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 10.10.0.1 (?), v2
    Info source: 10.10.0.1 (?), via bootstrap, priority 0, holdtime 150
         Uptime: 04:13:05, expires: 00:02:06
```

IP address of RP is known

We know the path to RP

```
R8#sh ip route 10.10.0.1
Routing entry for 10.10.0.1/32
  Known via "ospf 1", distance 110, metric 31, type intra area
  Last update from 10.0.38.3 on Ethernet0/2, 04:13:16 ago
  Routing Descriptor Blocks:
  * 10.0.38.3, from 10.10.0.1, 04:13:16 ago, via Ethernet0/2
      Route metric is 31, traffic share count is 1
```

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (3):
  - Router can't send PIM Join – incorrect path to RP

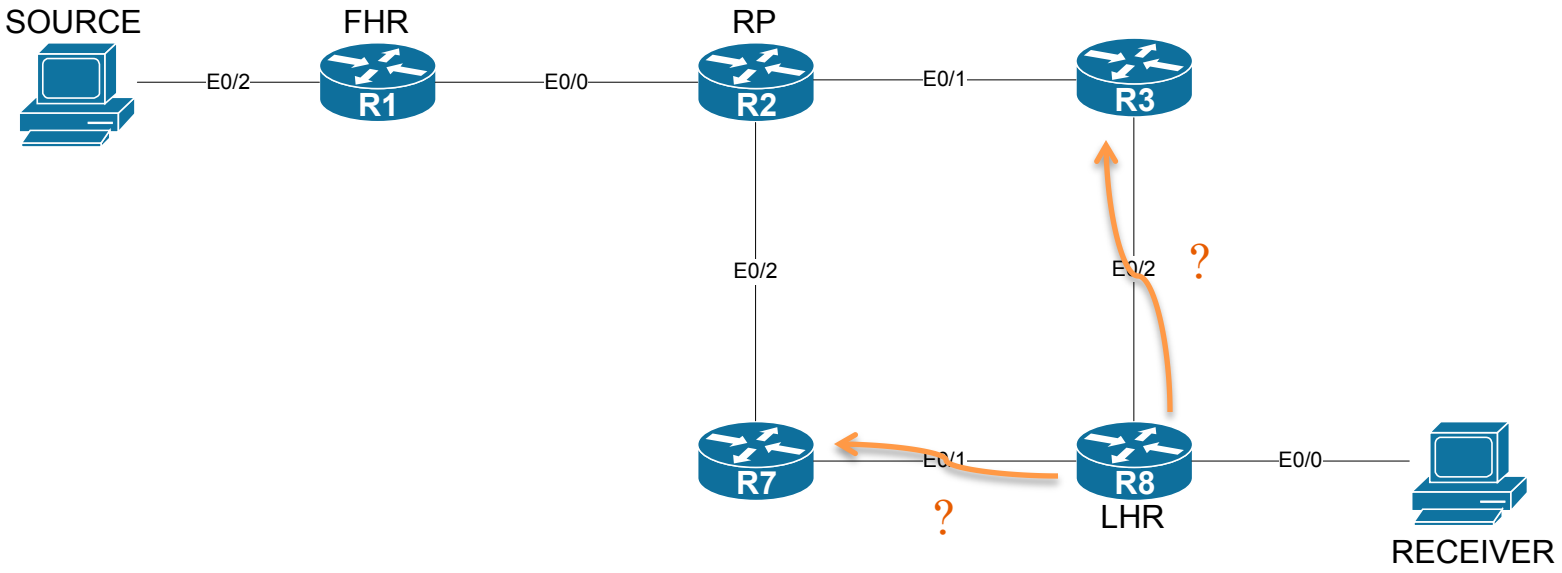

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (3):
  - Router can't send PIM Join – incorrect path to RP

```
R8#sh ip mroute 224.10.0.1

(*, 224.10.0.1), 2d04h/00:02:22, RP 10.10.0.2, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse, 2d04h/00:02:21

R8#show ip rpf 10.10.0.2
 failed, no route exists
```

PIM don't know where the source of traffic is

PIM don't know his neihghbor for this group

RPF rule is not fulfilled

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (3):
  - Router can't send PIM Join – incorrect path to RP
- Where is the problem?:
  - PIM not enabled on interface
  - Asymmetric routing
  - Load balancing

PROBLEM: RECEIVER CAN'T REGISTER

- Possible causes (3):
  - Router can't send PIM Join – incorrect path to RP

```
R8#sh ip mroute 224.10.0.1

(*, 224.10.0.1), 00:17:20/00:02:41, RP 10.10.0.2, flags: SJCL
  Incoming interface: Ethernet0/2, RPF nbr 10.0.38.3
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse, 00:17:20/00:02:41

R8#show ip rpf 10.10.0.2
RPF information for ? (10.10.0.2)
  RPF interface: Ethernet0/2
  RPF neighbor: ? (10.0.38.3)
  RPF route/mask: 10.10.0.2/32
  RPF type: unicast (ospf 1)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

PROBLEM: RECEIVER CAN'T REGISTER

- Problem solved :)
  - We have all information required to build PIM Join (*,G) tree

```
R8#sh ip mroute 224.10.0.1

(*, 224.10.0.1), 00:17:20/00:02:41, RP 10.10.0.2, flags: SJCL
  Incoming interface: Ethernet0/2, RPF nbr 10.0.38.3
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse, 00:17:20/00:02:41
```

PROBLEM: RECEIVER CAN'T REGISTER

- What if problem exists somewhere on path between LHR and RP?

- Just repeat diagnostic procedure on every router on this path

PROBLEM: SOURCE IS NOT REGISTERING AT RP

- **Problem 1:**
  - FHR router is not receiving nor processing multicast data

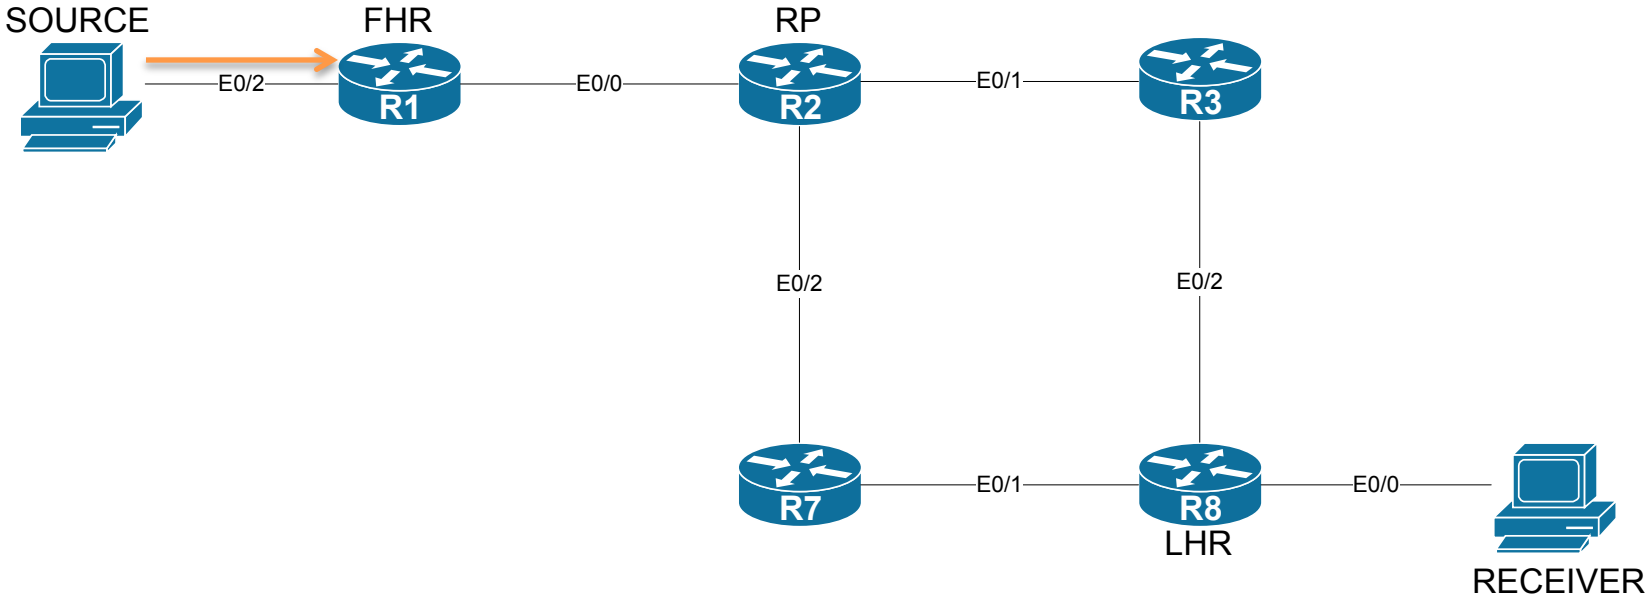

PROBLEM: SOURCE IS NOT REGISTERING AT RP

- **Problem 1:**
  - FHR router is not receiving nor processing multicast data

- Possible causes:
  - L2 IGMP filtering occurs on path between sender and the router
  - Access-list on router interface (another point of filtering!)
  - PIM is not enabled on interface

PROBLEM: SOURCE IS NOT REGISTERING AT RP

- **Problem 2:**
  - Source registration is not forwarded to RP

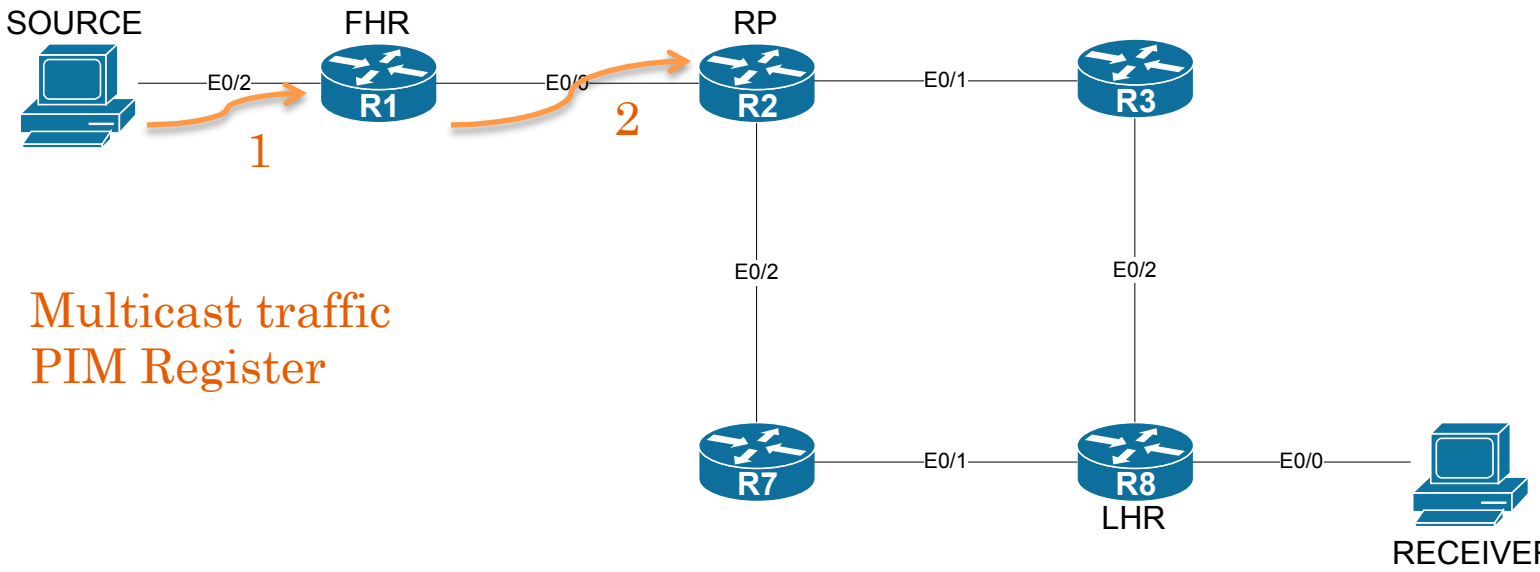

1. Multicast traffic
2. PIM Register

PROBLEM: SOURCE IS NOT REGISTERING AT RP

- Possible causes:
  - FHR router does not know RP for the group
  - FHR router does not know path to RP for the group
  - FHR router does not know correct neighbour on the path to the RP – incorrect path to RP

```
R1#show ip mroute

(*, 224.0.1.40), 01:07:32/00:02:32, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback0, Forward/Sparse-Dense, 01:07:32/00:00:00
    Ethernet0/0, Forward/Sparse-Dense, 01:07:31/00:00:00
```

PROBLEM: SOURCE IS NOT REGISTERING AT RP

- Possible causes:
  - FHR router does not know RP for the group
  - FHR router does not know path to RP for the group
  - FHR router does not know correct neighbour on the path to the RP – incorrect path to RP

```
R1#show ip mroute 224.10.0.1 count
IP Multicast Statistics
3 routes using 1520 bytes of memory
2 groups, 0.50 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 224.10.0.1, Source count: 1, Packets forwarded: 1, Packets received: 1
  RP-tree: Forwarding: 0/0/0/0, Other: 152/152/0
```

PROBLEM: SOURCE IS NOT REGISTERING AT RP

- In next step RP builds (S,G) tree to source of multicast traffic

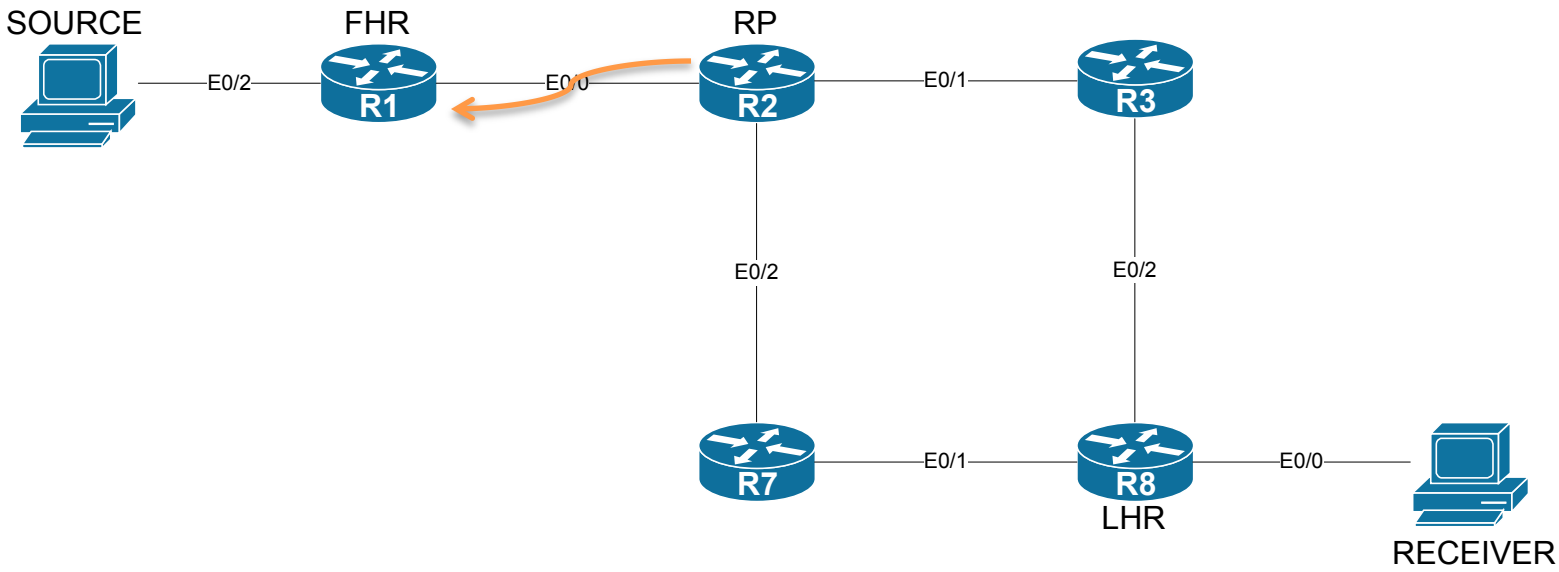

PROBLEM: SOURCE IS NOT REGISTERING AT RP

- In next step RP builds (S,G) tree to source of multicast traffic

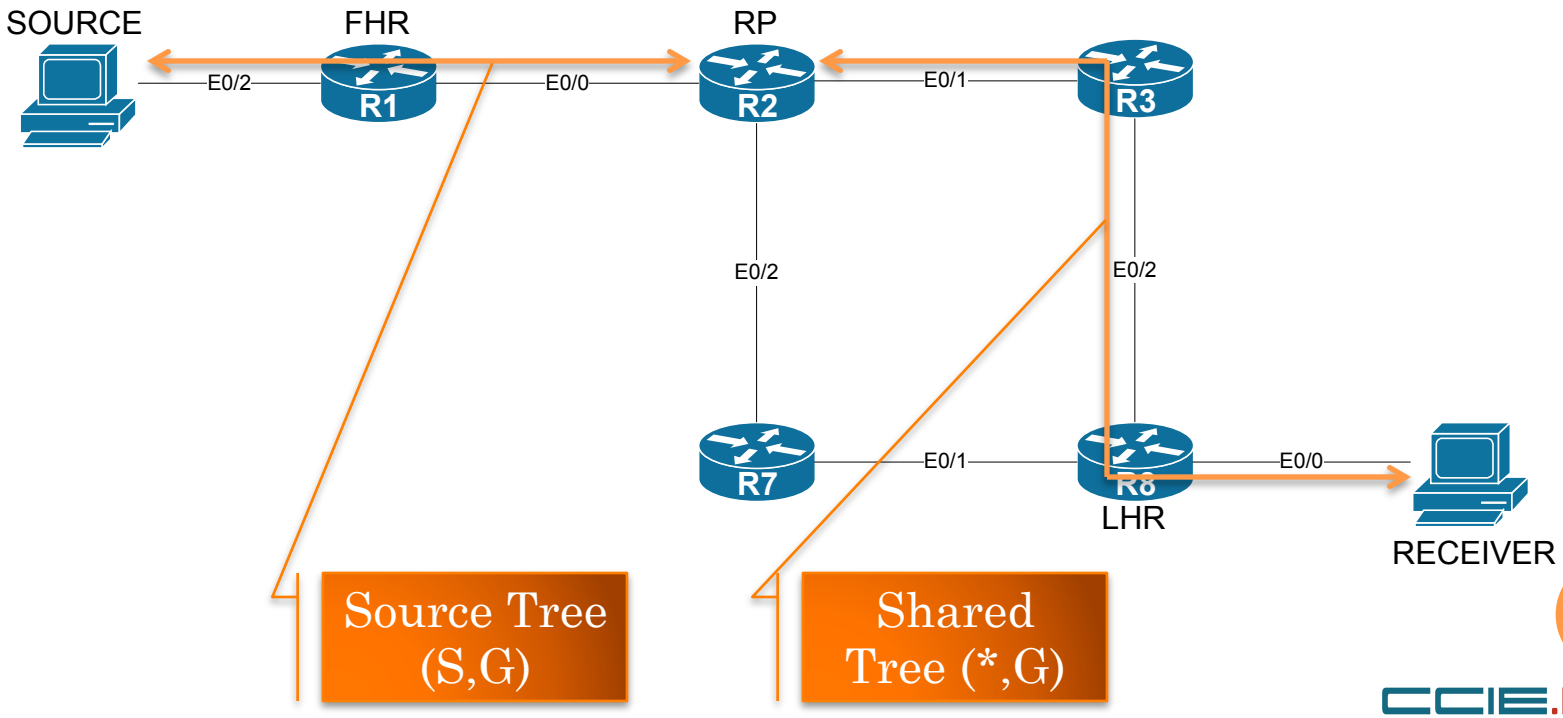

Summary

- Things, we have to check:
  - Is `ip multicast-routing` is present in configuration?
  - Is PIM enabled on all interfaces participating in multicast forwarding (including loopbacks)?
  - If we are using PIM-SM or PIM-BiDir check if RP is properly configured and it's IP address is know for every router within network (AutoRP, BSR, Static)
  - Are there any problems with unicast routing?
  - Is RPF requirement fulfilled?
  - Is security properly configured? (ACLs, policers, multicast boundary, BSR boundary, TTL etc.)

Questions?

CCIE.PL

POWERING your KNOWLEDGE.

Thank you